arXiv:1511.08152v1 [cs.DS] 25 Nov 2015

Max-Cut under Graph Constraints

Jon Lee*, Viswanath Nagarajan, and Xiangkun Shen

IOE Dept., University of Michigan, Ann Arbor, MI 48109, USA.

Abstract. An instance of the graph-constrained max-cut (GCMC) problem consists of (i) an undirected graph $G = (V, E)$ and (ii) edge-weights $c : \binom{V}{2} \to \mathbb{R}_+$ on a complete undirected graph. The objective is to find a subset $S \subseteq V$ of vertices satisfying some graph-based constraint in G that maximizes the weight $\sum_{u \in S, v \notin S} c_{uv}$ of edges in the cut $(S, V \setminus S)$. The types of graph constraints we can handle include independent set, vertex cover, dominating set and connectivity. Our main results are for the case when G is a graph with bounded treewidth, where we obtain a $\frac{1}{2}$-approximation algorithm. Our algorithm uses an LP relaxation based on the Sherali-Adams hierarchy. It can handle any graph constraint for which there is a (certain type of) dynamic program that exactly optimizes linear objectives.

Using known decomposition results, these imply essentially the same approximation ratio for GCMC under constraints such as independent set, dominating set and connectivity on a planar graph G (more generally for bounded-genus or excluded-minor graphs).

1 Introduction

The max-cut problem is an extensively studied combinatorial-optimization problem. Given an undirected edge-weighted graph, the goal is to find a subset $S \subseteq V$ of vertices that maximizes the weight of edges in the cut $(S, V \setminus S)$. Max-cut has a 0.878-approximation algorithm [14] which is known to be best-possible assuming the "unique games conjecture" [17]. It also has a number of practical applications, e.g., in circuit layout, statistical physics and clustering.

In some applications, one needs to solve the max-cut problem under additional constraints on the subset S. Consider for example, the following clustering problem. The input is an undirected graph $G = (V, E)$ representing, say, a social network (vertices V denote users and edges E denote connections between users), and a weight function $c : \binom{V}{2} \to \mathbb{R}_+$ representing, a dissimilarity measure between pairs of users. The goal is to find a subset $S \subseteq V$ of users that are connected in G while maximizing the weight of edges in the cut $(S, V \setminus S)$. This corresponds to finding a cluster of connected users that is as different as possible from its complement set. This "connected max-cut" problem also arises in image segmentation applications [22,16].

* Research of J. Lee was partially supported by NSF grant CMMI–1160915 and ONR grant N00014-14-1-0315.

Designing algorithms for constrained versions of max-cut is also interesting from a theoretical standpoint. For max-cut under certain types of constraints (such as cardinality or matroid constraints) good approximation algorithms are known, e.g., [2,1]. In fact, many of these results have since been extended to the more general setting of submodular objectives [12,9]. However, not much is known for max-cut under "graph-based" constraints as in the example above.

In this paper, we study a large class of graph-constrained max-cut problems and present unified approximation algorithms for them. Our results require that the constraint be defined on a graph G of bounded treewidth. (Treewidth is a measure of how similar a graph is to a tree structure — see §2 for definitions.) We note however that for a number of constraints (including the connectivity example above), we can combine our algorithm with known decomposition results [10,11] to obtain essentially the same approximation ratios when the constraint graph G is planar/bounded-genus/excluded-minor.

Problem definition. The input to the *graph-constrained max-cut* (GCMC) problem consists of (i) an n-vertex undirected graph $G = (V, E)$ which implicitly specifies a collection $\mathcal{C}_G \subseteq 2^V$ of feasible vertex subsets, and (ii) (symmetric) edge-weights $c : \binom{V}{2} \to \mathbb{R}_+$. The GCMC problem is then as follows:

$$\max_{S \in \mathcal{C}_G} \quad \sum_{u \in S, v \notin S} c(u, v). \tag{1}$$

In this paper, we assume that the constraint graph G has bounded treewidth. We also assume that the graph constraint $\mathcal{C}_G$ admits an exact dynamic program for optimizing a linear objective, i.e. for:

$$\max_{S \in \mathcal{C}_G} \quad \sum_{u \in S} f(u), \qquad \text{where } f : V \to \mathbb{R} \text{ is any given vertex weights.} \tag{2}$$

Note that the GCMC objective (1) is a quadratic function of the solution S, whereas our assumption (2) involves a *linear* function of the solution S. See §2 for more precise definitions/assumptions.

1.1 Our Results and Techniques

Our main result can be stated informally as follows.

Theorem 1 (GCMC result — informal). *Consider any instance of the* GCMC *problem on a bounded-treewidth graph $G = (V, E)$. Suppose there is an exact dynamic program for optimizing any linear function subject to constraint $\mathcal{C}_G$. Then we obtain a $\frac{1}{2}$-approximation algorithm for* GCMC*.*

This algorithm uses a linear-programming relaxation for GCMC based on the dynamic program (for linear objectives) which is further strengthened via the Sherali-Adams LP hierarchy. The resulting LP has polynomial size whenever the number of dynamic program states associated with a single tree-decomposition

node is constant (see §2 for the formal definition).[1] The rounding algorithm is a natural top-down procedure that randomly chooses a "state" for each tree-decomposition node using the LP's probability distribution conditional on the choices at its ancestor nodes. The final solution is obtained by combining the chosen states at each tree-decomposition node, which is guaranteed to satisfy constraint $\mathcal{C}_G$ due to properties of the dynamic program. We note that the choice of variables in the Sherali-Adams LP as well as the rounding algorithm are similar to those used in [15] for the sparsest cut problem on bounded-treewidth graphs. An important difference in our result is that we apply the Sherali-Adams hierarchy to a non-standard LP that is defined using the dynamic program for linear objectives. (If we were to apply Sherali-Adams to the standard LP, then it is unclear how to enforce the constraint $\mathcal{C}_G$ during the rounding algorithm.) Another difference is that our rounding algorithm needs to make a correlated choice in selecting the states of sibling nodes in order to satisfy constraint $\mathcal{C}_G$— this causes the number of variables in the Sherali-Adams LP to increase, but it still remains polynomial since the tree-decomposition has constant degree.

The requirement in Theorem 1 on the graph constraint $\mathcal{C}_G$ is satisfied by several interesting constraints and thus we obtain approximation algorithms for all these GCMC problems. See Section 4 for details.

Theorem 2 (Applications). *There is a $\frac{1}{2}$-approximation algorithm for* GCMC *under the following constraints in a bounded-treewidth graph: independent set, vertex cover, dominating set, connectivity.*

We note that many other constraints such as precedence, connected dominating set, and triangle matching also satisfy our requirement. In the interest of space, we only present details for the constraints mentioned in Theorem 2. We also note that for some of these constraints (e.g., independent set) one can come up with a problem specific algorithm where the approximation ratio depends on the treewidth k. Our result is stronger since the algorithm is more general, and the ratio is independent of k.

For many of the constraints above, we can use known decomposition results [10,11] to obtain approximation algorithms for GCMC when the constraint graph has bounded genus or excludes some fixed minor (e.g., planar graphs).

Corollary 1. *There is a $(\frac{1}{2} - \epsilon)$-approximation algorithm for* GCMC *under the following constraints in an excluded-minor graph: independent set, vertex cover, dominating set. Here $\epsilon > 0$ is a fixed constant.*

Corollary 2. *There is a $(\frac{1}{2} - \epsilon)$-approximation algorithm for connected max-cut in a bounded-genus graph. Here $\epsilon > 0$ is a fixed constant.*

Our approach can also handle other types of objectives. If $g : 2^V \to \mathbb{R}_+$ is the sum of a polynomial number of functions each of which is monotone, submodular and defined on a constant-size subset of V, then we obtain a $(1 - \frac{1}{e})$-approximation

[1] For other polynomial time dynamic programs, the LP has *quasi-polynomial* size.

algorithm for the problem of maximizing $g(S)$ subject to $S \in \mathcal{C}_G$. The graph constraint $\mathcal{C}_G$ is as above.[2] The main idea is to use the correlation gap of monotone submodular function.[3,9]

1.2 Related Work

For the basic undirected max-cut problem, there is an elegant 0.878-approximation algorithm [14] via semidefinite programming. This is also the best one can hope for, assuming the *unique games conjecture* [17].

Most of the prior work on constrained max-cut has focused on cardinality, matroid and knapsack constraints [2,1,12,9,18,19]. Constant-factor approximation algorithms are known for max-cut under the intersection of any constant number of such constraints — these results hold in the substantially more general setting of non-negative submodular functions. The main techniques used here are local search and the multilinear extension [8] of submodular functions. These results made crucial use of certain exchange properties of the underlying constraints, which are not true for graph-based constraints that we consider.

Closer to our setting, a version of the connected max-cut problem was studied recently in [16], where the connectivity constraint as well as the weight function were defined on the *same* graph G. The authors obtained an $O(\log n)$-approximation algorithm for general graphs, and an exact algorithm on bounded-treewidth graphs (which implied a PTAS for bounded-genus graphs); their algorithms relied heavily on the uniformity of the constraint/weight graphs. In contrast, we consider the connected max-cut problem where the connectivity constraint and the weight function are *unrelated*; in particular, our problem generalizes max-cut even when G is a trivial graph (e.g., a star). Moreover, our algorithms work for a much wider class of constraints. We note however that our results require graph G to have bounded treewidth — this is also necessary since some of the constraints we consider (e.g., independent set) are inapproximable in general graphs. (For connected max-cut itself, obtaining a non-trivial approximation ratio when G is a general graph remains an open question.)

In terms of techniques, the closest work to ours is [15]. We use ideas from [15] in formulating the (polynomial size) Sherali-Adams LP as well as in the rounding algorithm. There are important differences too, as discussed in §1.1.

Finally, our result adds to a somewhat small list [6,20,5,4,15,13] of algorithmic results based on the Sherali-Adams [21] LP hierarchy. We are not aware of a more direct approach to obtain a constant-factor approximation algorithm even for connected max-cut when the constraint graph G is a tree.

2 Preliminaries

Basic definitions. For an undirected complete graph on vertices V and subset $S \subseteq V$, let δS be the set of edges with exactly one end-point in S. For any weight function $c : \binom{V}{2} \to \mathbb{R}_+$ and subset $F \subseteq \binom{V}{2}$, we use $c(F) := \sum_{e \in F} c_e$.

[2] This setting is interesting only for constraints such as independent set, triangle matching and precedence that are not "upward closed".

Tree Decomposition. Given an undirected graph $G = (V, E)$, this consists of a tree $\mathcal{T} = (I, F)$ and a collection of vertex subsets $\{X_i \subseteq V\}_{i \in I}$ such that:

- for each $v \in V$, the nodes $\{i \in I : v \in X_i\}$ are connected in $\mathcal{T}$, and
- for each edge $(u, v) \in E$, there is some node $i \in I$ with $u, v \in X_i$.

The width of such a tree-decomposition is $\max_{i \in I}(|X_i| - 1)$, and the treewidth of G is the smallest width of any tree-decomposition for G.

We will work with "rooted" tree-decompositions that also specify a root node $r \in I$. The depth d of such a tree-decomposition is the length of the longest root-leaf path in $\mathcal{T}$. The depth of any node $i \in I$ is the length of the $r - i$ path in $\mathcal{T}$.

For any $i \in I$, the set V_i denotes all the vertices at or below node i, that is

$$V_i \quad := \quad \cup_{k \in \mathcal{T}_i} X_k, \quad \text{where } \mathcal{T}_i = \{k \in I : k \text{ in subtree of } \mathcal{T} \text{ rooted at } i\}.$$

The following known result provides a convenient representation of $\mathcal{T}$.

Theorem 3 (Balanced Tree Decomposition [7]). *Let $G = (V, E)$ be a graph of treewidth k. Then G has a rooted tree-decomposition $(\mathcal{T} = (I, F), \{X_i | i \in I\})$ where $\mathcal{T}$ is a binary tree of depth $2\lceil \log_{\frac{5}{4}}(2|V|) \rceil$ and treewidth at most $3k + 2$. This tree-decomposition can be found in $O(|V|)$ time.*

Dynamic program for linear objectives. We assume that the constraint $\mathcal{C}_G$ admits an exact dynamic programming (DP) algorithm for optimizing linear objectives, i.e. for the problem (2). There is some additional notation that is needed to formally describe the DP: this is necessary due to the generality of our results.

Definition 1 (DP) *With any tree-decomposition $(\mathcal{T} = (I, F), \{X_i | i \in I\})$, we associate the following:*

1. *For each node $i \in I$, there is a state space Σ_i.*
2. *For each node $i \in I$ and $\sigma \in \Sigma_i$, there is a collection $\mathcal{H}_{i,\sigma} \subseteq 2^{V_i}$ of subsets.*
3. *For each node $i \in I$, its children nodes $\{j, j'\}$ and $\sigma \in \Sigma_i$, there is a collection $\mathcal{F}_{i,\sigma} \subseteq \Sigma_j \times \Sigma_{j'}$ of valid combinations of children states.*

Assumption 1 (Linear objective DP for $\mathcal{C}_G$) *Let $(\mathcal{T} = (I, F), \{X_i | i \in I\})$ be any tree-decomposition. Then there exist Σ_i, $\mathcal{F}_{i,\sigma}$ and $\mathcal{H}_{i,\sigma}$ (see Definition 1) that satisfy the following:*

1. (bounded state space) *Σ_i and $\mathcal{F}_{i,\sigma}$ are all bounded by constant, that is, $\max_i |\Sigma_i| = t$ and $\max_{i,\sigma} |\mathcal{F}_{i,\sigma}| = p$, where $t, p = O(1)$.*
2. (required state) *For each $i \in I$ and $\sigma \in \Sigma_i$, the intersection with X_i of every set in $\mathcal{H}_{i,\sigma}$ is the same, denoted $X_{i,\sigma}$, that is $S \cap X_i = X_{i,\sigma}$ for all $S \in \mathcal{H}_{i,\sigma}$.*
3. (subproblem) *For each non-leaf node $i \in I$ with children $\{j, j'\}$ and $\sigma \in \Sigma_i$,*

$$\mathcal{H}_{i,\sigma} \quad = \quad \left\{ X_{i,\sigma} \cup S_j \cup S_{j'} \ : \ S_j \in \mathcal{H}_{j,w_j},\, S_{j'} \in \mathcal{H}_{j',w_{j'}},\, (w_j, w_{j'}) \in \mathcal{F}_{i,\sigma} \right\}.$$

By condition 2, for any leaf $\ell \in I$ and $\sigma \in \Sigma_\ell$, we have $\mathcal{H}_{\ell,\sigma} = \{X_{\ell,\sigma}\}$ or $\emptyset$.

4. (cover all constraints) *At the root node r, we have* $\mathcal{C}_G = \bigcup_{\sigma \in \Sigma_r} \mathcal{H}_{r,\sigma}$.

The most restrictive assumption is the first condition. To the best of our knowledge, all natural constraints that admit a dynamic program on bounded-treewidth graphs (for linear objectives) satisfy conditions 2-4. Even in cases when condition 1 is not true, a relaxed version holds (where t and p are polynomial), and our approach gives a *quasi-polynomial* time $\frac{1}{2}$-approximation algorithm.

Example: Here we outline how independent set satisfies the above requirements.

- The state space of each node $i \in I$ consists of all independent subsets of X_i.
- The subsets $\mathcal{H}_{i,\sigma}$ consist of all independent subsets $S \subseteq V_i$ with $S \cap X_i = \sigma$.
- The valid combinations $\mathcal{F}_{i,\sigma}$ consist of all tuples $(w_j, w_{j'})$ where the child states w_j and $w_{j'}$ are "consistent" with state σ at node i.

A formal proof of why the independent-set constraint satisfies Assumption 1 appears in Section 4. There, we also discuss a number of other graph constraints satisfying our assumption.

The following result follows from Assumption 1.

Claim 1 *For any $S \in \mathcal{C}_G$, there is a collection $\{b(i) \in \Sigma_i\}_{i \in I}$ such that:*

- *for each node $i \in I$ with children j and j', $(b(j), b(j')) \in \mathcal{F}_{i,b(i)}$,*
- *for each leaf ℓ we have $\mathcal{H}_{\ell,b(\ell)} \neq \emptyset$, and*
- $S = \bigcup_{i \in I} X_{i,b(i)}$.

Moreover, for any vertex $u \in V$, if $\overline{u} \in I$ denotes the highest node containing u then $u \in S \iff u \in X_{\overline{u},b(\overline{u})}$.

Proof. We define the states $b(i)$ in a top-down manner; we will also define an associated subset $B_i \in \mathcal{H}_{i,b(i)}$ at each node i. At the root, we set $b(r) = \sigma$ such that $S \in \mathcal{H}_{r,\sigma}$: this is well-defined by Assumption 1(4). We also set $B_r = S$. Having set $b(i)$ and $B_i \in \mathcal{H}_{i,b(i)}$ for any node $i \in I$ with children $\{j, j'\}$, we use Assumption 1(3) to write

$$B_i = X_{i,b(i)} \cup S_j \cup S_{j'} \quad \text{where } S_j \in \mathcal{H}_{j,w_j},\ S_{j'} \in \mathcal{H}_{j',w_{j'}} \text{ and } (w_j, w_{j'}) \in \mathcal{F}_{i,b(i)}.$$

Then we set $b(j) = w_j$ and $B_j = S_j$ for all the children J of node i. It is now easy to verify the first three conditions in the claim.

Since $S = \bigcup_{i \in I} X_{i,b(i)}$, it is clear that $u \in X_{\overline{u},b(\overline{u})} \implies u \in S$. In the other direction, suppose $u \notin X_{\overline{u},b(\overline{u})}$: we will show $u \notin S$. Since $\overline{u}$ is the highest node containing u, it suffices to show that $u \notin B_{\overline{u}}$ above. But this follows directly from Assumption 1(2) since $B_{\overline{u}} \in \mathcal{H}_{\overline{u},b(\overline{u})}$, $u \in X_{\overline{u}}$ and $u \notin X_{\overline{u},b(\overline{u})}$. ∎

Sherali-Adams LP hierarchy. This is one of the several "lift-and-project" procedures that, given a $\{0,1\}$ integer program, produces systematically a sequence of increasingly tighter convex relaxations. The Sherali-Adams procedure [21] involves generating stronger *LP relaxations* by adding new variables and constraints. The r^{th} round of this procedure has a variable $y(S)$ for every subset S of at most r variables in the original integer program — the new variable $y(S)$ corresponds to the joint event that all the original variables in S are one.

3 Approximation Algorithm for GCMC

In this section, we prove:

Theorem 4. *Consider any instance of the* GCMC *problem on a bounded-treewidth graph $G = (V, E)$. If the graph constraint $\mathcal{C}_G$ satisfies Assumption 1 then we obtain a*

$\frac{1}{2}$*-approximation algorithm.*

Algorithm outline: We start with a balanced tree-decomposition $\mathcal{T}$ of graph G, as given in Theorem 3; recall the associated definitions from §2. Then we formulate an LP relaxation of the problem using Assumption 1 (i.e. the dynamic program for linear objectives) and further strengthened by applying the Sherali-Adams operator. Finally we use a natural top-down rounding that relies on Assumption 1 and the Sherali-Adams constraints.

3.1 Linear Program

We start with some additional notation related to the tree-decomposition $\mathcal{T}$ (from Theorem 3) and our dynamic program assumption (Assumption 1).

- For any node $i \in I$, T_i is the set of nodes on the $r - i$ path along with the children of all nodes except i on this path. See also Figure 1.
- $\mathcal{P}$ is the collection of all node subsets J such that $J \subseteq T_{\ell_1} \cup T_{\ell_2}$ for some pair of leaf-nodes ℓ_1, ℓ_2. See also Figure 1.
- $s(i) \in \Sigma_i$ denotes a state at node i. Moreover, for any subset of nodes $N \subseteq I$, we use the shorthand $s(N) := \{s(k) : k \in N\}$.
- $\bar{u} \in I$ denotes the highest tree-decomposition node containing vertex u.

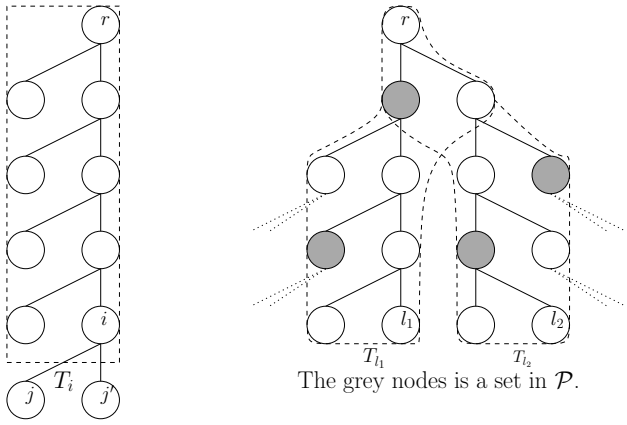

Fig. 1. Examples of (i) a set T_i and (ii) a set in $\mathcal{P}$.

The variables in our LP are $y(s(N))$ for all $\{s(k) \in \Sigma_k\}_{k \in N}$ and $N \in \mathcal{P}$. Variable $y(s(N))$ corresponds to the joint event that the solution (in $\mathcal{C}_G$) "induces" state $s(k)$ (in terms of Assumption 1) at each node $k \in N$.

We also use variables z_{uv} defined in constraint (3) that measure the probability of an edge (u, v) being cut. Constraints (4) are the Sherali-Adams constraints that enforce consistency among the y variables. Constraints (5)-(7) are from the dynamic program (Assumption 1) and require valid state selections.

$$\text{maximize} \sum_{\{u,v\} \in \binom{V}{2}} c_{uv} z_{uv} \tag{LP}$$

$$z_{uv} = \sum_{\substack{s(\bar{u}) \in \Sigma_{\bar{u}} \\ u \in X_{\bar{u}, s(\bar{u})}}} \sum_{\substack{s(\bar{v}) \in \Sigma_{\bar{v}} \\ v \notin X_{\bar{v}, s(\bar{v})}}} y(s(\{\bar{u}, \bar{v}\})) + \sum_{\substack{s(\bar{u}) \in \Sigma_{\bar{u}} \\ u \notin X_{\bar{u}, s(\bar{u})}}} \sum_{\substack{s(\bar{v}) \in \Sigma_{\bar{v}} \\ v \in X_{\bar{v}, s(\bar{v})}}} y(s(\{\bar{u}, \bar{v}\})),$$

$$\forall \{u, v\} \in \binom{V}{2}; \tag{3}$$

$$y(s(N)) = \sum_{s(i) \in \Sigma_i} y(s(N \cup \{i\})), \qquad \forall N \in \mathcal{P},\, i \notin N : N \cup \{i\} \in \mathcal{P}; \tag{4}$$

$$\sum_{s(r) \in \Sigma_r} y(s(r)) = 1; \tag{5}$$

$$y(s(\{i, j, j'\})) = 0, \qquad \forall i \in I,\, s(i) \in \Sigma_i,\, (s(j), s(j')) \notin \mathcal{F}_{i, s(i)}; \tag{6}$$

$$y(s(\ell)) = 0, \qquad \forall \ell \in I, s(\ell) \in \Sigma_\ell : \mathcal{H}_{\ell, s(\ell)} = \emptyset; \tag{7}$$

$$0 \le y(s(N)) \le 1, \qquad \forall N \in \mathcal{P},\, \{s(k) \in \Sigma_k\}_{k \in N}. \tag{8}$$

Claim 2 *For any node $i \in I$ with children j, j' and $s(k) \in \Sigma_k$ for all $k \in T_i$,*

$$y(s(T_i)) = \sum_{s(j) \in \Sigma_j} \sum_{s(j') \in \Sigma_{j'}} y(s(T_i \cup \{j, j'\})). \tag{9}$$

Proof. Note that $T_i \cup \{j, j'\} \subseteq T_\ell$ for any leaf node ℓ in the subtree below i. So $T_i \cup \{j, j'\} \in \mathcal{P}$ and the variables $y(s(T_i \cup \{j, j'\})$ are well-defined. The claim now follows by two applications of constraint (4). ∎

In constraint (6), we use j and j' to denote the two children of node $i \in I$.

3.2 The Rounding Algorithm

We start with the root node $r \in I$. Here $\{y(s(r)) : s(r) \in \Sigma_r\}$ defines a probability distribution over the states of r. We sample a state $a(r) \in \Sigma_r$ from this distribution. Then we continue top-down: given the chosen state $a(i)$ of any

node i, we sample states for both children of i *simultaneously* from their joint distribution given at node i.

Input : Optimal solution of LP.
Output: A vertex set in $\mathcal{C}_G$.

Sample a state $a(r)$ at the root node by distribution $y(s(r))$;
Do *process all nodes i in $\mathcal{T}$ in order of increasing depth* **:**
  Sample states $a(j), a(j')$ for the children of node i by joint distribution

$$\Pr[a(j) = s(j) \text{ and } a(j') = s(j')] \;=\; \frac{y(s(T_i \cup \{j, j'\}))}{y(s(T_i))}, \tag{10}$$

  where $s(T_i) = a(T_i)$.
end
Do *process all nodes i in $\mathcal{T}$ in order of decreasing depth* **:**
  $R_i = X_{i,a(i)} \cup R_j \cup R_{j'}$ where j, j' are the children of i.
end
$R = R_r$;
return R.

Algorithm 1: Rounding Algorithm for LP

3.3 Algorithm Analysis

Lemma 1. (LP) *is a valid relaxation of* GCMC.

Proof. Let $S \in \mathcal{C}_G$ be any feasible solution to the GCMC instance. Let $\{b(i)\}_{i \in I}$ denote the states given by Claim 1 corresponding to S. For any subset $N \in \mathcal{P}$ of nodes, and for all $\{s(i) \in \Sigma_i\}_{i \in N}$, set

$$y(s(N)) = \begin{cases} 1, & \text{if } s(i) = b(i) \text{ for all } i \in N; \\ 0, & \text{otherwise.} \end{cases}$$

It is easy to see that constraints (4) and (8) are satisfied. By the first two properties in Claim 1, it follows that constraints (6) and (7) are also satisfied. The last property in Claim 1 implies that $u \in S \iff u \in X_{\overline{u}, b(\overline{u})}$ for any vertex $u \in V$. So any edge $\{u, v\}$ is cut exactly when one of the following occurs:

- $u \in X_{\overline{u}, b(\overline{u})}$ and $v \notin X_{\overline{v}, b(\overline{v})}$;
- $u \notin X_{\overline{u}, b(\overline{u})}$ and $v \in X_{\overline{v}, b(\overline{v})}$.

Using the setting of variable z_{uv} in (3) it follows that z_{uv} is exactly the indicator of edge $\{u, v\}$ being cut by S. Thus the objective value in (LP) is $c(\delta S)$. ∎

Lemma 2. (LP) *has a polynomial number of variables and constraints. Hence the overall algorithm runs in polynomial time.*

Proof. There are $\binom{n}{2} = O(n^2)$ variables z_{uv}. Since the tree is binary, we have $|T_i| \leq 2d$ for any node i, where $d = O(\log n)$ is the depth of the tree-decomposition.

Moreover there are only $O(n^2)$ pairs of leaves as there are $O(n)$ leaf nodes. For each pair ℓ_1, ℓ_2 of leaves, we have $|T_{\ell_1} \cup T_{\ell_2}| \leq 4d$. Thus $|\mathcal{P}| \leq O(n^2) \cdot 2^{4d} = poly(n)$. By Assumption 1, we have $\max |\mathcal{H}_{i,\sigma}| = t = O(1)$, so the number of y-variables is at most $|\mathcal{P}| \cdot t^{4d} = poly(n)$. This shows that (LP) has polynomial size and can be solved optimally in polynomial time. Finally, it is easy to see that the rounding algorithm runs in polynomial time. ∎

Lemma 3. *The algorithm's solution R is always feasible.*

Proof. Note that the distribution used in Step 1 is always valid due to Claim 2; so the states $a(i)$s are well-defined.

We now show that for any node $i \in I$ with children j, j' we have $(a(j), a(j')) \in \mathcal{F}_{i,a(i)}$. Indeed, at the iteration for node i (when $a(j)$ and $a(j')$ are set) using the probability distribution in (10) and by constraint (6), we obtain that $(a(j), a(j')) \in \mathcal{F}_{i,a(i)}$ with probability one.

We show that for each node $i \in I$, the subset $R_i \in \mathcal{H}_{i,a(i)}$ by induction on the height of i. The base case is when i is a leaf. In this case, due to constraint (7) (and the validity of the rounding algorithm) we know that $\mathcal{H}_{i,a(i)} \neq \emptyset$. So $R_i = X_{i,a(i)} \in \mathcal{H}_{i,a(i)}$ by Assumption 1(3). For the inductive step, consider node $i \in I$ with children j, j' where $R_j \in \mathcal{H}_{j,a(j)}$ and $R_{j'} \in \mathcal{H}_{j',a(j')}$. Moreover, from the property above, $(a(j), a(j')) \in \mathcal{F}_{i,a(i)}$. Now using Assumption 1(3) we have $R_i = X_{i,a(i)} \cup R_j \cup R_{j'} \in \mathcal{H}_{i,a(i)}$. Thus the final solution $R \in \mathcal{C}_G$. ∎

Claim 3 *A vertex u is contained in solution R if and only if $u \in X_{\bar{u},a(\bar{u})}$.*

Proof. This proof is identical to that of the last property in Claim 1. ∎

In the rest of this section, we show that every edge (u, v) is cut by solution R with probability at least $z_{uv}/2$, which would prove the algorithm's approximation ratio. Lemma 4 handles the case when $\bar{u} \in T_{\bar{v}}$ (the case $\bar{v} \in T_{\bar{u}}$ is identical). And Lemma 5 handles the (harder) case when $\bar{u} \notin T_{\bar{v}}$ and $\bar{v} \notin T_{\bar{u}}$.

We first state some useful claims before proving the lemmas.

Observation 1 (see [15] for a similar use of this principle) *Let X, Y be two jointly distributed $\{0,1\}$ random variables. Then* $\Pr(X = 1)\Pr(Y = 0) + \Pr(X = 0)\Pr(Y = 1) \geq \frac{1}{2}[\Pr(X = 0, Y = 1) + \Pr(X = 1, Y = 0)]$.

Proof. Let $\Pr(X = 0, Y = 0) = x$, $\Pr(X = 0) = a$, $\Pr(Y = 0) = b$. The probability table is as below:

x	$a - x$	a
$b - x$	$1 + x - a - b$	$1 - a$
b	$1 - b$	

Then we want to show $a(1-b)+b(1-a) \geq \frac{1}{2}(a+b-2x) \Leftrightarrow a+b+2x \geq 4ab$. Since each probability is in $[0,1]$, we have $0 \leq x \leq \min\{a, b\}$ and $a + b - 1 \leq x$.

If $a+b>1$, we have $a+b+2x \geq 3a+3b-2>1$. If $ab<\frac{1}{4}$, it is done. If $\frac{1}{4} \leq ab \leq 1$, we have $6\sqrt{ab}-2 \geq 4ab$. Then $a+b+2x \geq 3a+3b-2 \geq 4ab$.

If $a+b \leq 1$, we have $a+b+2x \geq a+b \geq 2\sqrt{ab} \geq 4ab$.

Combine the above two cases, we have the observation is true. ∎

Claim 4 *For any node i and state $s(k) \in \Sigma_k$ for all $k \in T_i$, the rounding algorithm satisfies* $\Pr[a(T_i) = s(T_i)] = y(s(T_i))$.

Proof. We proceed by induction on the depth of node i. It is clearly true when $i = r$, i.e. $T_i = \{r\}$. Assuming the statement is true for node i, we will prove it for i's children. Let j, j' be the children nodes of i; note that $T_j = T_{j'} = T_i \cup \{j, j'\}$. Then using (10), we have

$$\Pr[a(T_j) = s(T_j) \mid a(T_i) = s(T_i)] \quad = \quad \frac{y(s(T_i \cup \{j, j'\}))}{y(s(T_i))}.$$

Combined with $\Pr[a(T_i) = s(T_i)] = y(s(T_i))$ we obtain $\Pr[a(T_j) = s(T_j)] = y(s(T_j))$ as desired. ∎

Claim 5 *For any $u, v \in V$, $s(\bar{u}) \in \Sigma_{\bar{u}}$ and $s(\bar{v}) \in \Sigma_{\bar{v}}$, we have*

$$y(s(\{\bar{u}, \bar{v}\})) = \sum_{\substack{s(k) \in \Sigma_k \\ k \in T_i \setminus \bar{u} \setminus \bar{v}}} y(s(T_i \cup \{\bar{u}, \bar{v}\})),$$

where i is the least common ancestor of $\bar{u}$ and $\bar{v}$.

Proof. Since i is the least common ancestor of $\bar{u}$ and $\bar{v}$, we have $T_i \cup \{\bar{u}, \bar{v}\} \in \mathcal{P}$. Then the claim follows by repeatedly applying constraint (4). ∎

Lemma 4. *Consider any $u, v \in V$ such that $\bar{u} \in T_{\bar{v}}$. Then the probability that edge (u, v) is cut by solution R is z_{uv}.*

Proof. Applying Claim 4 with node $i = \bar{v}$, for any $\{s(k) \in \Sigma_k : k \in T_{\bar{v}}\}$, we have $\Pr[a(T_{\bar{v}}) = s(T_{\bar{u}})] = y(s(T_{\bar{u}}))$. Let $D_u = \{s(\bar{u}) \in \Sigma_{\bar{u}} | u \in s(\bar{u})\}$ and $D_v = \{s(\bar{v}) \in \Sigma_{\bar{v}} | v \in s(\bar{v})\}$. Since $\bar{u} \in T_{\bar{v}}$,

$$\Pr[u \in R, v \notin R] = \sum_{s(\bar{u}) \in D_u} \sum_{s(\bar{v}) \notin D_v} \sum_{\substack{s(k) \in \Sigma_k \\ k \in T_{\bar{v}} \setminus \bar{u} \setminus \bar{v}}} y(s(T_{\bar{u}})) = \sum_{s(\bar{u}) \in D_u} \sum_{s(\bar{v}) \notin D_v} y(s(\bar{u}, \bar{v}))$$

The last equality above is by repeated application of constraint (4). Similarly we have

$$\Pr[u \notin R, v \in R] = \sum_{s(\bar{u}) \notin D_u} \sum_{s(\bar{v}) \in D_v} y(s(\bar{u}, \bar{v})),$$

which combined with constraint (5) implies $\Pr[|\{u, v\} \cap R| = 1] = z_{uv}$. ∎

Lemma 5. *Consider any $u, v \in V$ such that $\bar{u} \notin T_{\bar{v}}$ and $\bar{v} \notin T_{\bar{u}}$. Then the probability that edge (u, v) is cut by solution R is at least $z_{uv}/2$.*

Proof. In order to simplify notation, we define:

$$z^+_{uv} \;=\; \sum_{\substack{s(\bar{u})\in\Sigma_{\bar{u}}\\ u\in X_{\bar{u},s(\bar{u})}}} \sum_{\substack{s(\bar{v})\in\Sigma_{\bar{v}}\\ v\notin X_{\bar{v},s(\bar{v})}}} y(s(\{\bar{u},\bar{v}\})), \quad z^-_{uv} \;=\; \sum_{\substack{s(\bar{u})\in\Sigma_{\bar{u}}\\ u\notin X_{\bar{u},s(\bar{u})}}} \sum_{\substack{s(\bar{v})\in\Sigma_{\bar{v}}\\ v\in X_{\bar{v},s(\bar{v})}}} y(s(\{\bar{u},\bar{v}\})).$$

Note that $z_{uv} = z^+_{uv} + z^-_{uv}$.

Let $D_u = \{s(\bar{u}) \in \Sigma_{\bar{u}} | u \in s(\bar{u})\}$ and $D_v = \{s(\bar{v}) \in \Sigma_{\bar{v}} | v \in s(\bar{v})\}$. Let i denote the least common ancestor of nodes $\bar{u}$ and $\bar{v}$. For any choice of states $\{s(k) \in \Sigma_k\}_{k\in T_i}$ define:

$$z^+_{uv}(s(T_i)) = \sum_{s(\bar{u})\in D_u} \sum_{s(\bar{v})\notin D_v} \frac{y(s(T_i \cup \{\bar{u},\bar{v}\}))}{y(s(T_i))},$$

and similarly $z^-_{uv}(s(T_i))$.

In the rest of the proof we fix states $\{s(k) \in \Sigma_k\}_{k\in T_i}$ and condition on the event $\mathcal{E}$ that $a(T_i) = s(T_i)$. We will show:

$$\Pr[|\{u,v\} \cap R| = 1 \,|\, \mathcal{E}] \;\geq\; \frac{1}{2}\left(z^+_{uv}(s(T_i)) + z^-_{uv}(s(T_i))\right). \tag{11}$$

By taking expectation over the conditioning $s(T_i)$, this would imply Lemma 5.

We now define the following indicator random variables (conditioned on $\mathcal{E}$).

$$I_u = \begin{cases} 0 & \text{if } a(\bar{u}) \notin D_u \\ 1 & \text{if } a(\bar{u}) \in D_u \end{cases} \quad \text{and} \quad I_v = \begin{cases} 0 & \text{if } a(\bar{v}) \notin D_v \\ 1 & \text{if } a(\bar{v}) \in D_v \end{cases}.$$

Observe that I_u and I_v (*conditioned* on $\mathcal{E}$) are independent since $\bar{u} \notin T_{\bar{v}}$ and $\bar{v} \notin T_{\bar{u}}$. So,

$$\Pr[|\{u,v\} \cap R| = 1 \,|\, \mathcal{E}] \;=\; \Pr[I_u = 1]\cdot\Pr[I_v = 0] + \Pr[I_u = 0]\cdot\Pr[I_v = 1] \tag{12}$$

For any $s(k) \in \Sigma_k$ for $k \in T_{\bar{u}} \setminus T_i$, we have by Claim 4 and $T_i \subseteq T_{\bar{u}}$ that

$$\Pr[a(T_{\bar{u}}) = s(T_{\bar{u}}) \,|\, a(T_i) = s(T_i)] = \frac{\Pr[a(T_{\bar{u}}) = s(T_{\bar{u}})]}{\Pr[a(T_i) = s(T_i)]} = \frac{y(s(T_{\bar{u}}))}{y(s(T_i))}.$$

Therefore

$$\Pr[I_u = 1] \;=\; \sum_{s(\bar{u})\in D_u} \sum_{\substack{k\in T_{\bar{u}}\setminus T_i\setminus\{\bar{u}\}\\ s(k)\in\Sigma_k}} \frac{y(s(T_{\bar{u}}))}{y(s(T_i))} = \sum_{s(\bar{u})\in D_u} \frac{y(s(T_i \cup \{\bar{u}\}))}{y(s(T_i))}.$$

The last equality follows from the (4) constraint. Similarly,

$$\Pr[I_v = 1] \;=\; \sum_{s(\bar{v})\in D_v} \frac{y(s(T_i \cup \{\bar{v}\}))}{y(s(T_i))}.$$

Now define $\{0,1\}$ random variables X and Y jointly distributed as:

	$Y=0$	$Y=1$
$X=0$	$\Pr[I_v=1]-z^-_{uv}(s(T_i))$	$z^-_{uv}(s(T_i))$
$X=1$	$z^+_{uv}(s(T_i))$	$\Pr[I_u=1]-z^+_{uv}(s(T_i))$

Note that $\Pr[X=1]=\Pr[I_u=1]$ and $\Pr[Y=1]=\Pr[I_v=1]$. So, applying Observation 1 and using (12) we have:

$$\Pr[|\{u,v\}\cap R|=1\,|\,\mathcal{E}] \;\geq\; \frac{1}{2}\left(\Pr[X=0,Y=1]+\Pr[X=1,Y=0]\right),$$

which implies (11). ∎

4 Applications

In this section, we show a number of graph constraints that satisfy Assumption 1 and thereby obtain $\frac{1}{2}$-approximation algorithms for GCMC under these constraints (on bounded-treewidth graphs).

Recall that the underlying graph G is given by its tree-decomposition ($\mathcal{T}=(I,F),\{X_i|i\in I\}$) from Theorem 3. Recall also the definition of a dynamic program on this tree-decomposition, as given in Definition 1.

4.1 Independent Set

Given graph $G=(V,E)$ and edge-weights $c:\binom{V}{2}\to\mathbb{R}_+$ we want to maximize $c(\delta S)$ where S is an independent set in G.

For each node $i\in I$ define state space $\Sigma_i=\{\sigma\subseteq X_i\,|\,\sigma \text{ is an independent set}\}$. For each node $i\in I$ and $\sigma\in\Sigma_i$, we define:

- set $X_{i,\sigma}=\sigma$.
- collection $\mathcal{H}_{i,\sigma}=\{S\subseteq V_i\,|\,X_i\cap S=\sigma \text{ and } S \text{ is an independent set in } G[V_i]\}$.
- $\mathcal{F}_{i,\sigma}=\{(w_{j_1},w_{j_2})\,|\,\text{for each } j\in\{j_1,j_2\},\, w_j\in\Sigma_j \text{ such that } w_j\cap X_i=\sigma\cap X_j\}$ which denotes valid combinations. Note that the condition $w_j\cap X_i=\sigma\cap X_j$ enforces w_j to agree with σ on vertices of $X_i\cap X_j$.

We next show that these satisfy all the conditions in Assumption 1.

Assumption 1 part 1. We have $t=\max|\Sigma_i|\leq 2^k=O(1)$ for bounded-treewidth k. Also $p=\max|\mathcal{F}_{i,\sigma}|\leq t^2$ since each node has at most two children.

Assumption 1 part 2. By definition, for any $S\in\mathcal{H}_{i,\sigma}$ we have $S\cap X_i=\sigma=X_{i,\sigma}$.

Assumption 1 part 3. For any leaf $\ell\in I$ and $\sigma\in\Sigma_\ell$ it is clear that $\mathcal{H}_{\ell,\sigma}=\{X_{i,\sigma}\}$.

Consider now any non-leaf node i and $\sigma\in\Sigma_i$. Let

$$\mathcal{Z} \;=\; \{X_{i,\sigma}\cup S_{j_1}\cup S_{j_2}\,:\,S_{j_1}\in\mathcal{H}_{j_1,w_{j_1}},S_{j_2}\in\mathcal{H}_{j_2,w_{j_2}},(w_{j_1},w_{j_2})\in\mathcal{F}_{i,\sigma}\}. \quad (13)$$

We first prove $\mathcal{H}_{i,\sigma}\subseteq\mathcal{Z}$. For any $S\in\mathcal{H}_{i,\sigma}$ and child $j\in\{j_1,j_2\}$ let $S_j=S\cap V_j$ and $w_j=S\cap X_j$; since S is independent S_j is also an independent set,

and $S_j \in \mathcal{H}_{j,w_j}$. Note that $S \cap X_i = X_{i,\sigma}$. Since $V_i = X_i \cup V_{j_1} \cup V_{j_2}$, we have $S = X_{i,\sigma} \cup S_{j_1} \cup S_{j_2}$. Moreover, we have $\sigma \cap X_j = S \cap X_i \cap X_j = w_j \cap X_i$ for each $j \in \{j_1, j_2\}$. So we have $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$ and hence $S \in \mathcal{Z}$.

We next prove $\mathcal{Z} \subseteq \mathcal{H}_{i,\sigma}$. Consider any $S = X_{i,\sigma} \cup S_{j_1} \cup S_{j_2}$ as in (13). For $j \in \{j_1, j_2\}$ by definition of $\mathcal{F}_{i,\sigma}$ and $\mathcal{H}_{j,w_j}$, we have $\sigma \cap X_j = w_j \cap X_i = (S_j \cap X_j) \cap X_i$; since $X_i \cap (V_j \setminus X_j) = \emptyset$ (by definition of the tree-decomposition) we have $X_i \cap S_j = X_i \cap X_j \cap S_j = \sigma \cap X_j$. Thus we have $X_i \cap S = \sigma$. It just remains to prove that S is an independent set in $G[V_i]$. Since S_{j_1}, S_{j_2} and $X_{i,\sigma}$ are independent sets, if S were not independent then we must have an edge (u, v) where $u \in V_{j_1} \cup X_i$ and $v \in V_{j_2} \setminus X_i$ (or the symmetric case); this is not possible due to the tree-decomposition. So $S \in \mathcal{H}_{i,\sigma}$.

Assumption 1 part 4. This follows directly from the definition of $\mathcal{H}_{i,\sigma}$.

4.2 Connectivity

Given graph $G = (V, E)$ and edge-weights $c : \binom{V}{2} \to \mathbb{R}_+$ we want to maximize $c(\delta S)$ where S is a connected vertex-set in G.

For each node $i \in I$ define the state space

$$\Sigma_i = \{(B_i, P_i) \mid B_i \subseteq X_i,\ P_i \text{ is a partition of } B_i\}.$$

Here a state $\sigma = (B_i, P_i)$ specifies which subset B_i of the vertices (in X_i) are included in the solution and what is the connectivity pattern P_i among them.

For each node $i \in I$ and $\sigma = (B_i, P_i) \in \Sigma_i$, we define:

- set $X_{i,\sigma} = B_i$.
- if $i = r$ (at the root) $\Sigma_r = \{(B_r, P_r) \mid B_r \subseteq X_r,\ P_r = \{B_r\}\}$.
- if $i \neq r$ then $\mathcal{H}_{i,\sigma} = \{S \subseteq V_i \mid X_i \cap S = B_i$, each part of P_i is connected in $G[S]$ and every connected component of $G[S]$ contains some vertex of $B_i\}$.
- partition $\bar{P}_i$ denotes the connected components in $G[B_i]$.
- $\mathcal{F}_{i,\sigma}$ consists of (w_{j_1}, w_{j_2}) where for $j \in \{j_1, j_2\}$, $w_j = (B_j, P_j) \in \Sigma_j$ such that $B_i \cap X_j = B_j \cap X_i$ and each part of P_j contains some vertex of B_i, and P_i is satisfied[3] by $\bar{P}_i \cup P_{j_1} \cup P_{j_2}$. Note that for some states there may be no such pair (w_{j_1}, w_{j_2}) : in this case $\mathcal{F}_{i,\sigma}$ is empty.

Assumption 1 part 1. For each node i, the possible number of vertex subsets B_i is at most 2^k and the possible number of partitions P_i is at most k^k, where k is the treewidth. So for a bounded-treewidth k, we have $t = \max |\Sigma_i| \leq k^{k+1} = O(1)$. Then $p = \max |\mathcal{F}_{i,\sigma}| \leq t^2 = O(1)$.

Assumption 1 part 2. This follows directly from the definition of $\mathcal{H}_{i,\sigma}$ and $X_{i,\sigma}$.

[3] Given two partitions Q and R, their union $P = Q \cup R$ is the refined partition where a pair of elements are in the same part iff they occur in the same part of either Q or R. Moreover, a partition P is said to be satisfied by another partition P' if P' is a refinement of P, i.e. every pair of elements in the same part of P also lie in the same part of P'.

Assumption 1 part 3. Let $\mathcal{Z}$ be as in (13) with the new definitions of $\mathcal{H}$ and $\mathcal{F}$ for connectivity (as above). The leaf case is trivial, so we consider a non-leaf node $i \in I$ and $\sigma = (B_i, P_i) \in \Sigma_i$. To reduce notation we just use j to denote a child of i; we will not specify $j \in \{j_1, j_2\}$ each time.

We first prove $\mathcal{H}_{i,\sigma} \subseteq \mathcal{Z}$. For any $S \in \mathcal{H}_{i,\sigma}$, let $S_j = V_j \cap S$ and $B_j = X_j \cap S$. Let P_j be a partition of B_j with a part $C \cap B_j$ for every connected component C in $G[S_j]$. Let $w_j = (B_j, P_j)$. We will show that $S_j \in \mathcal{H}_{j,w_j}$ and $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$.

- $S_j \in \mathcal{H}_{j,w_j}$. By definition of B_j, we have $X_j \cap S_j = X_j \cap V_j \cap S = X_j \cap S = B_j$. We only need to prove each connected component of $G[S_j]$ has at least one vertex of B_j. We will in fact show that each component of $G[S_j]$ has at least one vertex of B_i (i.e. in $B_j \cap B_i$). Suppose (for contradiction) there is some connected component C in $G[S_j]$ which does not have any vertex of B_i. By $S \in \mathcal{H}_{i,\sigma}$ we know that in the (larger) graph $G[S]$ component C has to be connected to some vertex $u \in B_i$. Then there is a path π in $G[S]$ from some vertex $u' \in C$ to u such that u' is the *only* vertex of C on π (see also Figure 2). Let (u', v') be the first edge of π, so $u' \in C \subseteq S_j$ and $v' \in S \setminus S_j$. By tree-decomposition, there is some node containing both u' and v'. Since $u', v' \in V_i$ and $u' \in S_j, v' \notin S_j$, that node can only be i. This means $u' \in B_i$, contrary to our assumption. Therefore we have $S_j \in \mathcal{H}_{j,w_j}$.
- $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$. We have $B_i \cap X_j = B_j \cap X_i$ by definition of w_j. Since we already proved that each connected component of $G[S_j]$ has at least one vertex of $B_j \cap B_i$, we know that each part of partition P_j has at least one vertex of B_i. By tree-decomposition we have $G[V_i] = G[X_i] \cup G[V_{j_1}] \cup G[V_{j_2}]$, so $G[S] = G[B_i] \cup G[S_{j_1}] \cup G[S_{j_2}]$. Hence partition P_i is satisfied by $\bar{P}_i \cup P_{j_1} \cup P_{j_2}$. Thus $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$.

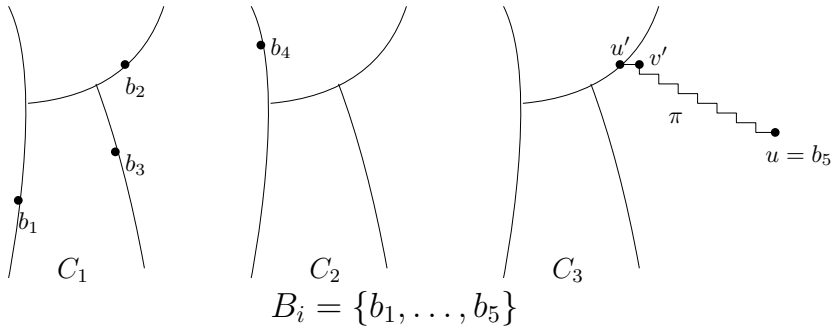

Fig. 2. Maximal connected component of $G[S_j]$

Next we prove $\mathcal{Z} \subseteq \mathcal{H}_{i,\sigma}$. Consider any $S \in \mathcal{Z}$ given by $S = B_i \cup S_{j_1} \cup S_{j_2}$ as in (13). The fact that $S \cap X_i = B_i$ follows exactly as in the case of an independent-set constraint. Since P_i is satisfied by $\bar{P}_i \cup P_{j_1} \cup P_{j_2}$ and S_j connects up each

part of P_j, it follows that $G[S] = G[B_i] \cup G[S_{j_1}] \cup G[S_{j_2}]$ connects up each part of P_i. It remains to show that each connected component of $G[S]$ has a vertex of B_i. Since $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$ we know that each part of P_j has a B_i-vertex. By $S_j \in \mathcal{H}_{j,w_j}$, we know that each component of $G[S_j]$ contains some vertex $u \in B_j$, and this vertex u is connected to some vertex $v \in B_i$ (as each part of P_j contains a B_i-vertex); so every component of $G[S_j]$ contains some vertex of B_i. Hence each component of $G[S] = G[B_i] \cup G[S_{j_1}] \cup G[S_{j_2}]$ also contains some vertex of B_i.

Assumption 1 part 4. By our definition of Σ_r, any solution given by $\mathcal{H}_{r,\sigma}$ requires all chosen vertices to be connected. Thus this assumption is satisfied.

4.3 Vertex Cover

Given graph $G = (V, E)$ and edge-weights $c : \binom{V}{2} \to \mathbb{R}_+$ we want to maximize $c(\delta S)$ where S is a vertex cover in G (i.e. S contains at least one end-point of each edge in E).

For each node $i \in I$ define the state space $\Sigma_i = \{\sigma \subseteq X_i \,|\, \sigma \text{ is a vertex cover in } G[X_i]\}$. For each node $i \in I$ and $\sigma \in \Sigma_i$, we define:

- set $X_{i,\sigma} = \sigma$.
- collection $\mathcal{H}_{i,\sigma} = \{S \subseteq V_i \,|\, X_i \cap S = \sigma \text{ and } S \text{ is a vertex cover in } G[V_i]\}$.
- $\mathcal{F}_{i,\sigma} = \{(w_{j_1}, w_{j_2}) \,|\, \text{for each } j \in \{j_1, j_2\},\ w_j \in \Sigma_j \text{ such that } w_j \cap X_i = \sigma \cap X_j\}$ which denotes valid combinations. Note that the condition $w_j \cap X_i = \sigma \cap X_j$ enforces w_j to agree with σ on vertices of $X_i \cap X_j$.

The proof of the above notation satisfying Assumption 1 is identical to the independent set proof.

4.4 Dominating Set

Given graph $G = (V, E)$ and edge-weights $c : \binom{V}{2} \to \mathbb{R}_+$ we want to maximize $c(\delta S)$ where S is a dominating set in G (i.e. every vertex in V is either in S or a neighbor of some vertex in S).

For each node $i \in I$ define the state space

$$\Sigma_i = \{(B_i, Y_i) \,|\, B_i \subseteq X_i,\ Y_i \subseteq X_i\}.$$

For vertex set $S \subseteq V$, we use $N(S)$ to denote S and the neighbor vertices of S.

Here a state $\sigma = (B_i, Y_i)$ specifies which subset B_i of the vertices (in X_i) are included in the solution and what subset Y_i of the vertices (in X_i) should be dominated.

For each node $i \in I$ and $\sigma = (B_i, Y_i) \in \Sigma_i$, we define:

- set $X_{i,\sigma} = B_i$
- if $i = r$ (at the root) $\Sigma_r = \{(B_r, Y_r) \,|\, B_r \subseteq X_r,\ Y_r = \emptyset\}$.
- if $i \neq r$ then $\mathcal{H}_{i,\sigma} = \{S \subseteq V_i |\ X_i \cap S = B_i, S \text{ is a dominating set of } V_i \setminus Y_i \text{ in } G[V_i]\}$

– $\mathcal{F}_{i,\sigma}$ consists of (w_{j_1}, w_{j_2}) where for $j \in \{j_1, j_2\}$, $w_j = (B_j, Y_j) \in \Sigma_j$ such that $B_i \cap X_j = B_j \cap X_i$ and $V_i \setminus Y_i \subseteq (V_{j_1} \setminus Y_{j_1}) \cup (V_{j_2} \setminus Y_{j_2}) \cup N(B_i)$. Note that for some states there may be no such pair (w_{j_1}, w_{j_2}) : in this case $\mathcal{F}_{i,\sigma}$ is empty.

Assumption 1 part 1. For each node i, the possible number of vertex subsets X_i, Y_i is at most 2^k, where k is the treewidth. So for a bounded-treewidth k, we have $t = \max |\Sigma_i| \leq 2^{2k} = O(1)$. Then $p = \max |\mathcal{F}_{i,\sigma}| \leq t^2 = O(1)$.

Assumption 1 part 2. This follows directly from the definition of $\mathcal{H}_{i,\sigma}$ and $X_{i,\sigma}$.

Assumption 1 part 3. Let $\mathcal{Z}$ be as in (13) with the new definitions of $\mathcal{H}$ and $\mathcal{F}$ for dominate set (as above). The leaf case is trivial, so we consider a non-leaf node $i \in I$ and $\sigma = (B_i, Y_i) \in \Sigma_i$. To reduce notation we just use j to denote a child of i; we will not specify $j \in \{j_1, j_2\}$ each time.

We first prove $\mathcal{H}_{i,\sigma} \subseteq \mathcal{Z}$. For any $S \in \mathcal{H}_{i,\sigma}$, let $S_j = V_j \cap S$ and $B_j = X_j \cap S$. Let $Y_j = X_j \setminus N(S_j)$. Let $w_j = (B_j, Y_j)$. We will show that $S_j \in \mathcal{H}_{j,w_j}$ and $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$.

– $S_j \in \mathcal{H}_{j,w_j}$. By definition of B_j, we have $X_j \cap S_j = X_j \cap V_j \cap S = X_j \cap S = B_j$. We only need to prove S_j is a dominating set of $V_j \setminus Y_j$ in $G[V_j]$. For all $v \in V_j \setminus Y_j$: If $v \in X_i$, since $v \in V_j$, we have $v \in X_j$. Since $v \notin Y_j$ and $v \in X_j$, by $Y_j = X_j \setminus N(S_j)$, we have $v \in N(S_j)$ by tree-decomposition, that is v is dominated by S_j. If $v \notin X_i$, we have $v \in V_i \setminus Y_i$. Then v is dominated by S. There is some $u \in S$ such that $(u, v) \in E$. Then by tree-decomposition, since $v \in V_j$ and $v \notin X_i$, we have $u \in V_j$. Then since $S_j = S \cap V_j$, we have $u \in S_j$. v is dominated by S_j. Then we have S_j will dominate $V_j \setminus Y_j$. Thus we have $S_j \in \mathcal{H}_{j,w_j}$.
– $(w_{j_1}, w_{j_2}) \in \mathcal{F}_{i,\sigma}$. We have $B_i \cap X_j = B_j \cap X_i$ and $Y_j = X_j \cap (Y_i \cup N(B_j))$ by definition of w_j. It remains to show that $V_i \setminus Y_i \subseteq (V_{j_1} \setminus Y_{j_1}) \cup (V_{j_2} \setminus Y_{j_2}) \cup N(B_i)$. For all $v \in V_i \setminus Y_i$, we have v is dominated by S. There is $u \in S$ such that $(u, v) \in E$. If $u \in X_i$, then we have $u \in B_i$ thus $v \in N(B_i)$. If $u \in V_j \setminus X_i$, we have $u \in S_j$. By tree-decomposition, $u \notin X_i$, $u \in V_j$ and $(u, v) \in E$ gives us $v \in V_j$. If $v \notin X_j$, we have $v \in V_j \setminus Y_j$. If $v \in X_j$, since $u \in S_j$, $v \in N(u)$, we have $v \in N(S_j)$. Then by $Y_j = X_j \setminus N(S_j)$, we have $v \notin Y_j$. Thus $v \in V_j \setminus Y_j$. Therefore, for all $v \in V_i \setminus Y_i$, we have $v \in (V_{j_1} \setminus Y_{j_1}) \cup (V_{j_2} \setminus Y_{j_2}) \cup N(B_i)$. Thus $V_i \setminus Y_i \subseteq (V_{j_1} \setminus Y_{j_1}) \cup (V_{j_2} \setminus Y_{j_2}) \cup N(B_i)$.

Next we prove $\mathcal{Z} \subseteq \mathcal{H}_{i,\sigma}$. Consider any $S \in \mathcal{Z}$ given by $S = B_i \cup S_{j_1} \cup S_{j_2}$ as in (13). The fact that $S \cap X_i = B_i$ follows exactly as in the case of an independent-set constraint. It remains to show that S is a dominating set of $V_i \setminus Y_i$. For all $v \in V_i \setminus Y_i$, we have $v \in (V_{j_1} \setminus Y_{j_1}) \cup (V_{j_2} \setminus Y_{j_2}) \cup N(B_i)$. Since S_j is a dominate set of $V_j \setminus Y_j$ and B_i is a dominate set of $N(B_i)$, we have v is dominated by $B_i \cup S_{j_1} \cup S_{j_2}$, v is dominated by S. Thus we have S is a dominating set of $V_i \setminus Y_i$. $S \in \mathcal{H}_{i,\sigma}$.

Assumption 1 part 4. By our definition of Σ_r, any solution given by $\mathcal{H}_{r,\sigma}$ requires all vertices are dominated. Thus this assumption is satisfied.

5 Bounded-genus and Excluded-minor Graphs

Here we use known decomposition results to show that our results can be extended to a larger class of graphs, and prove Corollary 1 and 2.

5.1 Excluded-minor graph

Recall the following decomposition of any excluded-minor graph into graphs of bounded treewidth.

Theorem 5. *[11] For a fixed graph H, there exists a constant c_H such that, for any integer $h \geq 1$ and for every H-minor-free graph G, the vertices of G can be partitioned into $h+1$ sets such that any h of that sets induce a graph of treewidth at most $c_H h$. Furthermore, such partition can be found in polynomial time.*

Algorithm 2 for Corollary 1 is given below. For a subset $V_i \subseteq V$, let G_i be the graph obtained by contracting V_i to v_{new}. Then the edge weight on G_i is defined as

$$c_i(u,v) = \begin{cases} c(u,v), \text{ if } u,v \in V \setminus V_i \\ \sum_{w \in V_i} c(u,w), \text{ if } u \in V \setminus V_i, v = v_{new} \end{cases} \tag{14}$$

We have v_{new} can increase treewidth by at most one since we can add it to each tree node and give a feasible tree-decomposition.

We will show Corollary 1 with the following claims.

Claim 6 *Let $V_1, \ldots, V_h$ be a partition of V. Let S' be any vertex subset of V. Then there is some i such that $c(\delta(S' \setminus V_i)) \geq (1 - \frac{2}{h})c(\delta S')$.*

Proof. Since $V_i, \ldots, V_h$ is a partition of V, we have $\sum_{i=1}^{h} c(\delta(S' \cap V_i)) \leq 2c(\delta S')$. Then $\min_i c(\delta(S' \cap V_i)) \leq \frac{2}{h} c(\delta S') \Leftrightarrow \max_i c(\delta(S' \setminus V_i)) \geq (1 - \frac{2}{h})c(\delta S')$. ∎

Claim 7 *Let $V_1, \ldots, V_h$ be a partition of V. Let S' be any vertex subset of V. Then there is some i such that $c(\delta(S' \cup V_i)) \geq (1 - \frac{2}{h})c(\delta S')$.*

Proof. Since $V_i, \ldots, V_h$ is a partition of V, we have $\sum_{i=1}^{h} (c(\delta S') - c(\delta(S' \cup V_i))) \leq 2c(\delta S')$. Then $\min_i (c(\delta S') - c(\delta(S' \cup V_i))) \leq \frac{2}{h} c(\delta S') \Leftrightarrow \max_i c(\delta(S' \cup V_i)) \geq (1 - \frac{2}{h})c(\delta S')$. ∎

Claim 8 *Let V_i be a subset of V. Suppose S' is a feasible solution of some* GCMC *problem and $S \setminus V_i$ is a feasible solution in $G[V \setminus V_i]$ with same graph constraint. Then the solution S_i given by* GCMC *algorithm in $G[V \setminus V_i]$ has cut value $c(\delta S_i) \geq \frac{1}{2} c(\delta S' \setminus V_i)$.*

Proof. Since $S \setminus V_i$ is a feasible solution in $G[V \setminus V_i]$, we have the optimal solution, S^* in $G[V \setminus V_i]$ has cut value $c(\delta S^*) \geq c(\delta S' \setminus V_i)$. By Theorem 4, we have $c(\delta S_i) \geq \frac{1}{2} c(\delta S^*)$, thus we have $c(\delta S_i) \geq \frac{1}{2} c(\delta S' \setminus V_i)$. ∎

Input : H-minor-free graph G
Output: A vertex set in $\mathcal{C}_G$
Use Theorem 5 to partition V into $V_1, \ldots, V_h$;
if *independent-set constraint* **then**
  for $i = 1$ *to* h **do**
    Solve GCMC in $G[V \setminus V_i]$ with edge-weight c;
    Let S_i be the solution.;
    $S = \arg\max_{S \in \{S_i\}} c(\delta S_i)$
  end
end
if *vertex-cover or dominating-set constraint* **then**
  for $i = 1$ *to* h **do**
    Solve GCMC in G_i with edge-weight c_i and require v_{new} to be part of the solution;
    `/* this requirement can be achieved by adding constraints`
      $y(s(r)) = 0$ `for all` $v_{new} \notin s(r)$ `to` (LP) `*/`
    Let S_i' be the solution.;
    $S_i = S_i' \setminus \{v_{new}\} \cup V_i$;
    $i = \arg\max_{j=1\ldots h} c(\delta(S_j))$;
    $S = S_i$;
  end
end
return S;

Algorithm 2: Algorithm for excluded minor graph

Claim 9 *Let V_i be a subset of V. Let G_i and c_i be defined as* (14)*. Suppose S' is a feasible solution of some* GCMC *problem and $S' \setminus V_i \cup \{v_{new}\}$ is a feasible solution in G_i with same graph constraint. Then the solution S_i given by* GCMC *algorithm in G_i has cut value $c_i(\delta S_i) \geq \frac{1}{2} c(\delta(S' \cup V_i))$ and $c(\delta(S_i \setminus \{v_{new}\} \cup V_i)) = c_i(\delta S_i)$.*

Proof. Since $S' \setminus V_i \cup \{v_{new}\}$ is a feasible solution in G_i, we have the optimal solution, S^* in G_i has cut value $c_i(\delta S^*) \geq c_i(\delta(S' \setminus V_i \cup \{v_{new}\}))$. By Theorem 4, we have $c_i(\delta S_i) \geq \frac{1}{2} c_i(\delta S^*)$. Then by definition of c_i we have $c_i(\delta S_i) \geq \frac{1}{2} c(\delta(S' \cup V_i\}))$ and $c(\delta(S_i \setminus \{v_{new}\} \cup V_i)) = c_i(\delta S_i)$. ∎

Claim 10 *Let S^* be the optimal solution of* GCMC *with independent-set constraint. The solution S given by the Algorithm 2 is a feasible solution to the original* GCMC *problem and $c(\delta S) \geq \frac{1}{2}(1 - \frac{2}{h}) c(\delta S^*)$.*

Proof. Since S is an independent set in $G[V \setminus V_i]$, then it is an independent set in G. S is a feasible solution to the original GCMC problem. Also we have that $S^* \setminus V_i$ is a feasible solution in $G[V \setminus V_i]$. Then by Claim 8, we have $c(\delta S) \geq \frac{1}{2} c(\delta(S^* \setminus V_i))$ and by Claim 6, we have $c(\delta(S^* \setminus V_i)) \geq (1 - \frac{2}{h}) c(\delta S^*)$. Combine the last two inequalities, we have $c(\delta S) \geq \frac{1}{2}(1 - \frac{2}{h}) c(\delta S^*)$. ∎

Claim 11 *Let S^* be the optimal solution of* GCMC *with vertex-cover constraint. The solution S given by the Algorithm 2 is a feasible solution to the original* GCMC *problem and $c(\delta S) \geq \frac{1}{2}(1 - \frac{6}{h}) c(\delta S^*)$.*

Proof. Since S'_i is a vertex cover of G_i. Then by definition of G_i, $S_i \cup V_i$ is a vertex cover of G. And by definition of G_i and c_i, we have $S^* \cup \{v_{new}\} \setminus V_i$ is a feasible solution for G_i. Then by Claim 9, $c(\delta S) \geq \frac{1}{2}c(\delta S^* \cup V_i)$. And by Claim 7, we have $c(\delta S^* \cup V_i) \geq (1 - \frac{2}{h})c(\delta S^*)$. Combine the last two inequalities, we have $c(\delta S) \geq \frac{1}{2}(1 - \frac{2}{h})c(\delta S^*)$. ∎

Claim 12 *Let S^* be the optimal solution of* GCMC *with dominating-set constraint. The solution S given by the Algorithm 2 is a feasible solution to the original* GCMC *problem and $c(\delta S) \geq \frac{1}{2}(1 - \frac{2}{h})c(\delta S^*)$.*

Proof. Since S'_i is a dominating set of G_i. Then by definition of G_i, $S_i \cup V_i$ is a dominating set of G. And by definition of G_i and c_i, we have $S^* \cup \{v_{new}\} \setminus V_i$ is a feasible solution for G_i. Then by Claim 9, $c(\delta S) \geq \frac{1}{2}c(\delta S^* \cup V_i)$. And by Claim 7, we have $c(\delta S^* \cup V_i) \geq (1 - \frac{2}{h})c(\delta S^*)$. Combine the last two inequalities, we have $c(\delta S) \geq \frac{1}{2}(1 - \frac{2}{h})c(\delta S^*)$. ∎

5.2 Bounded-genus graph

Here we use:

Theorem 6. *[10] For a bounded-genus graph G and an integer h, the edge of G can be partitioned in h color classes $E_1, \ldots, E_h$ such that contracting all the edges in any color class leads to a graph with treewidth $O(h)$. Further, the color classes are obtained by a radial coloring and have the following property: If edge $e = (u, v)$ is in class i, then every edge e' such that $e \cap e' \neq \emptyset$ is in class $i - 1$ or i or $i + 1$.*

The proof of Corollary 2 using Theorem 6 is identical to the proof in [16] for the "uniform" connected max-cut problem. For each E_i, let S_i be the solution in G with E_i contracted, then $S'_i = \{v | v \in S_i$ or v is contracted to some vertex of $S_i\}$ is connected in G. Although our edge-weights c are defined on a complete graph, the proof of [16] still works. The main reason is that each vertex gets contracted in at most 3 of the graphs G with E_i contracted.

References

1. Ageev, A.A., Hassin, R., Sviridenko, M.: A 0.5-approximation algorithm for MAX DICUT with given sizes of parts. SIAM J. Discrete Math. 14(2), 246–255 (2001)
2. Ageev, A.A., Sviridenko, M.: Approximation algorithms for maximum coverage and max cut with given sizes of parts. In: IPCO. pp. 17–30 (1999)
3. Agrawal, S., Ding, Y., Saberi, A., Ye, Y.: Price of Correlations in Stochastic Optimization. Operations Research 60(1), 150–162 (2012)
4. Bansal, N., Lee, K.W., Nagarajan, V., Zafer, M.: Minimum congestion mapping in a cloud. SIAM J. on Comput. 44(3), 819–843 (2015)
5. Bateni, M., Charikar, M., Guruswami, V.: Maxmin allocation via degree lower-bounded arborescences. In: STOC. pp. 543–552 (2009)
6. Bienstock, D., Özbay, N.: Tree-width and the Sherali-Adams operator. Discrete Optimization 1(1), 13–21 (2004)

7. Bodlaender, H.L.: NC-algorithms for graphs with small treewidth. In: Graph-theoretic concepts in computer science (Amsterdam, 1988), pp. 1–10 (1989)
8. Călinescu, G., Chekuri, C., Pál, M., Vondrák, J.: Maximizing a monotone submodular function subject to a matroid constraint. SIAM J. Comp. 40(6), 1740–66 (2011)
9. Chekuri, C., Vondrák, J., Zenklusen, R.: Submodular function maximization via the multilinear relaxation and contention resolution schemes. SIAM J. Comput. 43(6), 1831–1879 (2014)
10. Demaine, E.D., Hajiaghayi, M.T., Kawarabayashi, K.: Algorithmic graph minor theory: Decomposition, approximation, and coloring. In: FOCS. pp. 637–646 (2005)
11. Demaine, E.D., Hajiaghayi, M., Kawarabayashi, K.: Contraction decomposition in h-minor-free graphs and algorithmic applications. In: STOC. pp. 441–450 (2011)
12. Feldman, M., Naor, J., Schwartz, R.: A unified continuous greedy algorithm for submodular maximization. In: FOCS. pp. 570–579 (2011)
13. Friggstad, Z., Könemann, J., Kun-Ko, Y., Louis, A., Shadravan, M., Tulsiani, M.: Linear programming hierarchies suffice for directed steiner tree. In: IPCO. pp. 285–296 (2014)
14. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. J. Assoc. Comput. Mach. 42(6), 1115–1145 (1995)
15. Gupta, A., Talwar, K., Witmer, D.: Sparsest cut on bounded treewidth graphs: algorithms and hardness results. In: STOC. pp. 281–290 (2013)
16. Hajiaghayi, M., Kortsarz, G., MacDavid, R., Purohit, M., Sarpatwar, K.: Approximation algorithms for connected maximum cut and related problems. In: ESA. pp. 693–704 (2015)
17. Khot, S., Kindler, G., Mossel, E., O'Donnell, R.: Optimal inapproximability results for MAX-CUT and other 2-variable csps? SIAM J. Comput. 37(1), 319–357 (2007)
18. Lee, J., Mirrokni, V.S., Nagarajan, V., Sviridenko, M.: Maximizing nonmonotone submodular functions under matroid or knapsack constraints. SIAM J. Discrete Math. 23(4), 2053–2078 (2010)
19. Lee, J., Sviridenko, M., Vondrák, J.: Submodular maximization over multiple matroids via generalized exchange properties. Math. Op. Res. 35(4), 795–806 (2010)
20. Magen, A., Moharrami, M.: Robust algorithms for on minor-free graphs based on the sherali-adams hierarchy. In: APPROX. pp. 258–271 (2009)
21. Sherali, H.D., Adams, W.P.: A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. SIAM J. Discrete Math. 3(3), 411–430 (1990)
22. Vicente, S., Kolmogorov, V., Rother, C.: Graph cut based image segmentation with connectivity priors. In: CVPR (2008)